Globally Optimal Base Station Clustering in Interference Alignment-Based Multicell Networks

Rasmus Brandt, *Student Member, IEEE,* Rami Mochaourab, *Member, IEEE,*
and Mats Bengtsson, *Senior Member, IEEE*

arXiv:1602.08273v1 [cs.IT] 26 Feb 2016

***Abstract*—Coordinated precoding based on interference alignment is a promising technique for improving the throughputs in future wireless multicell networks. In small networks, all base stations can typically jointly coordinate their precoding. In large networks however, base station clustering is necessary due to the otherwise overwhelmingly high channel state information (CSI) acquisition overhead. In this work, we provide a branch and bound algorithm for finding the globally optimal base station clustering. The algorithm is mainly intended for benchmarking existing suboptimal clustering schemes. We propose a general model for the user throughputs, which only depends on the long-term CSI statistics. The model assumes intracluster interference alignment and is able to account for the CSI acquisition overhead. By enumerating a search tree using a best-first search and pruning sub-trees in which the optimal solution provably cannot be, the proposed method converges to the optimal solution. The pruning is done using specifically derived bounds, which exploit some assumed structure in the throughput model. It is empirically shown that the proposed method has an average complexity which is orders of magnitude lower than that of exhaustive search.**

I. Introduction

For coordinated precoding [1] in intermediate to large sized multicell networks, base station clustering [2]–[4] is necessary for reasons including channel state information (CSI) acquisition overhead, backhaul delays and implementation complexity constraints. In frequency-division duplex mode, the CSI acquisition overhead is due to the feedback required [5], [6], whereas in time-division duplex mode, the CSI acquisition overhead is due to pilot contamination and allocation [7], [8].

For the case of interference alignment (IA) precoding [9], suboptimal base station clustering algorithms have earlier been proposed in [2] where clusters are orthogonalized and a heuristic algorithm for the grouping was proposed, in [3] where the clusters are non-orthogonal and a heuristic algorithm on an interference graph was proposed, and in [4] where coalition formation and game theory was applied to a generalized frame structure. To the best of the authors' knowledge however, no works in the literature have addressed the problem of finding the globally optimal base station clustering for IA-based systems. Naive exhaustive search over all possible clusterings is not tractable, due to its super-exponential complexity. Yet, the globally optimal base station clustering is important in order to benchmark the more practical schemes in e.g [2]–[4]. Therefore, in this paper, we propose a structured method based on branch and bound [10], [11] for finding the globally optimal base station clustering. We consider a generalized throughput model which encompasses the models in [2]–[4]. When evaluated using the throughput model of [4], empirical evidence shows that the resulting algorithm finds the global optimum at an average complexity which is orders of magnitude lower than that of exhaustive search.

The authors are with the Department of Signal Processing, ACCESS Linnæus Centre, School of Electrical Engineering, KTH Royal Institute of Technology, Stockholm, Sweden. E-mails: rabr5411@kth.se, ramimo@kth.se, mats.bengtsson@ee.kth.se.

II. Problem Formulation

We consider a symmetric multicell network where I base stations (BSs) each serve K mobile stations (MSs) in the downlink. A BS together with its served MSs is called a *cell* and we denote the kth served MS by BS i as i_k. The BSs each have M antennas and the MSs each have N antennas. Each MS is served d spatial data streams. BS i allocates[1] a power of P_{i_k} to MS i_k, in total using a power of $P_i = \sum_{k=1}^{K} P_{i_k}$, and MS i_k has a thermal noise power of $\sigma_{i_k}^2$. The average large scale fading between BS j and MS i_k is $\gamma_{i_k j}$.

The cooperation between the BSs is determined by the BS clustering, which mathematically is described as a set partition:

Definition 1 (Set partition)**.** *A set partition $\mathcal{S} = \{\mathcal{C}_1, \ldots, \mathcal{C}_S\}$ is a partition of $\mathcal{I} = \{1, \ldots, I\}$ into disjoint and non-empty sets called clusters, such that $\mathcal{C}_s \subseteq \mathcal{I}$ for all $\mathcal{C}_s \in \mathcal{S}$ and $\bigcup_{s=1}^{S} \mathcal{C}_s = \mathcal{I}$. For a cell $i \in \mathcal{C}_s$, we let $\mathcal{S}(i) = \mathcal{C}_s$.*

We assume that IA is used to completely cancel the interference within each cluster.[2] Thus only the intercluster interference remains, which is reflected in the long-term signal-to-interference-and-noise ratios (SINRs) of the MSs:

Assumption 1 (Signal-to-interference-and-noise ratio)**.** *Let $\rho_{i_k} : 2^{\mathcal{I}} \to \mathbb{R}_+$ be the long-term SINR of MS i_k defined as*

$$\rho_{i_k}(\mathcal{S}(i)) = \frac{\gamma_{i_k i} P_{i_k}}{\sigma_{i_k}^2 + \sum_{j \in \mathcal{I} \setminus \mathcal{S}(i)} \gamma_{i_k j} P_j}. \tag{1}$$

We consider a general model for the MS throughputs, which depends on the cluster size and the long-term SINR. The cluster size determines the overhead, whereas the long-term SINR determines the achievable rate.

Assumption 2 (Throughput)**.** *For a cluster size $|\mathcal{S}(i)|$ and a long-term SINR $\rho_{i_k}(\mathcal{S}(i))$, the throughput of MS i_k is given by $t_{i_k}(\mathcal{S}) = v_{i_k}(|\mathcal{S}(i)|, \rho_{i_k}(\mathcal{S}(i)))$, where $v_{i_k} : \mathbb{N} \times \mathbb{R}_+ \to \mathbb{R}_+$ is unimodal in its first argument and non-decreasing in its second argument.*

[1] Any *fixed* power levels can be used, e.g. obtained from some single-cell power allocation method [1, Ch. 1.2]. Generalizing to *adaptive* multicell power allocation would however lead to loss of tractability in the SINR bound of Thm. 1, due to ρ_{i_k} not being supermodular [12] when the powers are adaptive.

[2] Within each cluster, both intra-cell and inter-cell interference is cancelled.

The structure of $\rho_{i_k}(\cdot)$ and the monotonicity properties of $v_{i_k}(\cdot,\cdot)$ will be used in the throughput bound to be derived below. The model in Assumption 2 is quite general and is compatible with several existing throughput models:

Example 1. *In [2], the clusters are orthogonalized using time sharing, and no intercluster interference is thus received. A coherence time of L_c is available. Each BS owns $1/I$ of the coherence time, which is contributed to the corresponding cluster. Larger clusters give more time for data transmission but also require more CSI feedback, which in [2] is modelled as a quadratic function, giving the throughput model as:*

$$v_{i_k}(|\mathcal{S}(i)|,\cdot) = \left(\frac{|\mathcal{S}(i)|}{I} - \frac{|\mathcal{S}(i)|^2}{L_c}\right) d\log\left(1+\varrho_{i_k}\right). \quad (2)$$

where $\varrho_{i_k} = \frac{\gamma_{i_k i} P_{i_k}}{\sigma^2_{i_k}}$ is the constant signal-to-noise ratio *(SNR). The function in* (2) *is strictly unimodal in its first argument and independent of its second argument.*

Example 2. *In [3], the clusters are operating using spectrum sharing. The CSI acquisition overhead is not accounted for. A slightly modified*[3] *version of their throughput model is then:*

$$v_{i_k}(\cdot,\rho_{i_k}(\mathcal{S}(i))) = d\log\left(1+\rho_{i_k}(\mathcal{S}(i))\right). \quad (3)$$

The function in (3) *is independent of its first argument, and strictly increasing in its second argument.*

Example 3. *In [4], intercluster time sharing and intercluster spectrum sharing are used in two different orthogonal phases. For the CSI acquisition overhead model during the time sharing phase, a model similar to the one in [2] is used. For the achievable rates during the spectrum sharing phase, long-term averages are derived involving an exponential integral. The model is thus*

$$v_{i_k}(|\mathcal{S}(i)|,\rho_{i_k}(\mathcal{S}(i))) = \alpha^{(1)}_{i_k}(|\mathcal{S}(i)|)\, r^{(1)}_{i_k} + r^{(2)}_{i_k}(\rho_{i_k}(\mathcal{S}(i))) \quad (4)$$

where

$$\alpha^{(1)}_{i_k}(|\mathcal{S}(i)|) = \frac{|\mathcal{S}(i)|}{I} - \frac{(M+K(N+d))\,|\mathcal{S}(i)| + KM\,|\mathcal{S}(i)|^2}{L_c},$$

$$r^{(1)}_{i_k} = d\, e^{1/\varrho_{i_k}} \int_{1/\varrho_{i_k}}^{\infty} t^{-1} e^{-t}\,\mathrm{d}t,$$

$$r^{(2)}_{i_k}(\rho_{i_k}(\mathcal{S}(i))) = d\, e^{1/\rho_{i_k}(\mathcal{S}(i))} \int_{1/\rho_{i_k}(\mathcal{S}(i))}^{\infty} t^{-1} e^{-t}\,\mathrm{d}t.$$

The function in (4) *is strictly unimodal in its first argument and strictly increasing in its second argument.*

Given the MS throughput model, we introduce the notion of a system-level objective:

Definition 2 (Objective)**.** *The performance of the entire multi-cell system is given by $f(\mathcal{S}) = g(t_{1_1}(\mathcal{S}),\ldots,t_{I_K}(\mathcal{S}))$, where $g\colon \mathbb{R}_+^{I\cdot K} \to \mathbb{R}_+$ is an argument-wise non-decreasing function.*

[3]The original SINR model in [3] includes the impact of the instantaneous IA filters, which we neglect here in order to avoid the cross-dependence between the IA solution and the clustering. This corresponds to how the approximated interference graph weights are derived in [3].

The function $f(\mathcal{S})$ thus maps a set partition to the corresponding system-level objective. Typical examples of objective functions are the weighted sum $f_{\text{WSR}}(\mathcal{S}) = \sum_{(i,k)} \lambda_{i_k} t_{i_k}(\mathcal{S})$ and the minimum weighted throughput $f_{\min}(\mathcal{S}) = \min_{(i,k)} \lambda_{i_k} t_{i_k}(\mathcal{S})$.

III. Globally Optimal Base Station Clustering

We will now provide a method for solving the following combinatorial optimization problem:

$$\begin{aligned} \mathcal{S}^\star = \arg\max_{\mathcal{S}} \quad & f(\mathcal{S}) \\ \text{subject to} \quad & \mathcal{S} \text{ satisfying Def. 1} \\ & |\mathcal{S}(i)| \le D,\ \forall\, i \in \mathcal{I}. \end{aligned} \quad (5)$$

The cardinality constraint is used to model cluster size constraints due to IA feasibility [13], CSI acquisition feasibility [4], implementation feasibility, etc.

A. Restricted Growth Strings and Exhaustive Search

In the algorithm to be proposed, we use the following alternate representation of a set partition:

Definition 3 (Restricted growth string, [14, Sec. 7.2.1.5])**.** *A set partition $\mathcal{S}$ can equivalently be expressed using a* restricted growth string *$a = a_1 a_2 \ldots a_I$ with the property that $a_i \le 1 + \max(a_1,\ldots,a_{i-1})$ for $i \in \mathcal{I}$. Then $a_i \in \mathbb{N}$ describes which cluster that cell i belongs to. We let $\mathcal{S}_a$ denote the mapping from a to the set partition $\mathcal{S}$, and $a_{\mathcal{S}}$ as its inverse.*

For example, the set partition $\mathcal{S}_a = \{\{1,3\},\{2\},\{4\}\}$ would be encoded as $a_{\mathcal{S}} = 1213$. One approach to solving the optimization problem in (5) is now by enumerating all restricted growth strings of length I, using e.g. Alg. H of [14, Sec. 7.2.1.5]. The complexity of this approach is however $\mathbb{B}_I$, the Ith *Bell number*[4], which grows super-exponentially.

B. Branch and Bound Algorithm

Most of the possible set partitions are typically not interesting in the sense of the objective of (5). For example, most set partitions will include clusters whose members are placed far apart, thus leading to low SINRs. By prioritizing set partitions with a potential to achieve large throughputs, the complexity of finding the globally optimal set partition can be decreased significantly compared to that of exhaustive search. This is the idea of the branch and bound approach [10], which entails bounding the optimal value $f(\mathcal{S}^\star)$ from above and below for a sequence of *partial solutions*. When the bounds converge, the optimal solution has been found. The partial solutions are described using *partial restricted growth strings*:

Definition 4 (Partial restricted growth string)**.** *The restricted growth string $\bar{a} = \bar{a}_1\bar{a}_2\ldots\bar{a}_l$ is* partial *if $l = \text{LENGTH}(\bar{a}) \le I$. The corresponding partial set partition, where only the first l cells are constrained into clusters, is denoted $\mathcal{S}_{\bar{a}}$.*

[4]The Ith Bell number describes the number of set partitions of $\mathcal{I}$ [15, p. 287], and can be bounded as $\mathbb{B}_I < (0.792I/\log(1+I))^I$ [16]. The 17 first Bell numbers are 1, 1, 2, 5, 15, 52, 203, 877, 4 140, 21 147, 115 975, 678 570, 4 213 597, 27 644 437, 190 899 322, 1 382 958 545, 10 480 142 147.

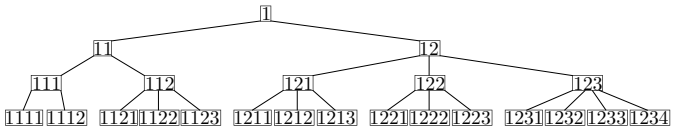

Fig. 1. Example of branch and bound search tree for $I = 4$.

The branch and bound method considers the sequence of partial solutions by dynamically exploring a search tree (see Fig. 1, at the top of the page), in which each interior node corresponds to a partial restricted growth string. By starting at the root and traversing down the search tree[5], more cells are constrained into clusters, ultimately giving the leaves of the tree which describe all possible restricted growth strings.

1) Bounds: We now provide the bounds that will be used to avoid exploring large parts of the search tree.

Lemma 1 (Objective bound)**.** *Let $\check{t}_{i_k}(\mathcal{S}_{\bar{a}})$ be an upper bound of the throughput of MS i_k for all leaf nodes in the sub-tree below the node described by $\bar{a}$. Then the function*

$$\check{f}(\mathcal{S}_{\bar{a}}) \;=\; g(\check{t}_{1_1}(\mathcal{S}_{\bar{a}}), \ldots, \check{t}_{I_K}(\mathcal{S}_{\bar{a}}))$$

is an upper bound of the objective in (5) *for all leaves in the sub-tree below the node described by $\bar{a}$.*

Proof: This follows directly from the argument-wise monotonicity of $g(t_{1_1}, \ldots, t_{I_K})$ in Def. 2. ∎

In order to describe the throughput bound, we will introduce three sets. Given a node described by $\bar{a}$, the cells in $\mathcal{P}_{\bar{a}} = \{1, \ldots, \text{LENGTH}(\bar{a})\} \subseteq \mathcal{I}$ are constrained into clusters as given by $\mathcal{S}_{\bar{a}}$. The remaining cells in $\mathcal{P}_{\bar{a}}^{\perp} = \mathcal{I} \setminus \mathcal{P}_{\bar{a}}$ are still unconstrained.[6] The set of cells which could accommodate more members in the corresponding clusters[7] are written as $\mathcal{F}_{\bar{a}} = \{i \in \mathcal{P}_{\bar{a}} : |\mathcal{S}_{\bar{a}}(i)| < D\} \cup \mathcal{P}_{\bar{a}}^{\perp}$.

Theorem 1 (Throughput bound)**.** *Let $t_{i_k}(\mathcal{S}_a)$ be the throughput of MS i_k for some leaf node in the sub-tree below the node described by $\bar{a}$. It can be bounded as $t_{i_k}(\mathcal{S}_a) \;=\; v_{i_k}(|\mathcal{S}_a(i)|, \rho_{i_k}(\mathcal{S}_a(i))) \leq v_{i_k}(\check{B}_{i_k}, \check{\rho}_{i_k})$ where*

$$\check{B}_{i_k} = \begin{cases} |\mathcal{S}_{\bar{a}}(i)| & \text{if } |\mathcal{S}_{\bar{a}}(i)| \geq B_{i_k}^{\star}, \\ \min\left(|\mathcal{S}_{\bar{a}}(i)| + \left|\mathcal{P}_{\bar{a}}^{\perp}\right|, B_{i_k}^{\star}\right) & \text{else if } i \in \mathcal{P}_{\bar{a}}, \\ \min\left(|\mathcal{F}_{\bar{a}}|, B_{i_k}^{\star}\right) & \text{else if } i \in \mathcal{P}_{\bar{a}}^{\perp}, \end{cases}$$

$$B_{i_k}^{\star} = \underset{b \in \mathbb{N}, b \leq D}{\arg\max}\, t_{i_k}(b, \check{\rho}_{i_k}),$$

and

$$\begin{aligned} \check{\rho}_{i_k} = \underset{|\mathcal{E}_{i_k}| \leq D}{\text{maximize}} \quad & \rho_{i_k}(\mathcal{E}_{i_k}) \\ \text{subject to} \quad & \text{if } i \in \mathcal{P}_{\bar{a}}\text{:} \\ & \qquad \mathcal{S}_{\bar{a}}(i) \subseteq \mathcal{E}_{i_k} \subseteq \left(\mathcal{S}_{\bar{a}}(i) \cup \mathcal{P}_{\bar{a}}^{\perp}\right) \\ & \text{else if } i \in \mathcal{P}_{\bar{a}}^{\perp}\text{:} \\ & \qquad \mathcal{E}_{i_k} \subseteq \mathcal{F}_{\bar{a}}. \end{aligned} \tag{6}$$

[5]At level $i \leq I$ of the tree, there are $\mathbb{B}_i$ nodes.

[6]In the sub-tree below the node described by $\bar{a}$, there is a leaf node for all possible ways of constraining the cells in $\mathcal{P}_{\bar{a}}^{\perp}$ into clusters.

[7]For the sake of this definition, we consider the non-constrained cells in $\mathcal{P}_{\bar{a}}^{\perp}$ to be in singleton clusters.

Algorithm 1 Branch and Bound for Base Station Clustering

Input: Initial $a_{\text{incumbent}}$ from some heuristic, $\epsilon \geq 0$

1: `live` $\leftarrow [1]$
2: **while** LENGTH(`live`) > 0 **do**
3: $\bar{a}_{\text{parent}} \leftarrow$ node from `live` with highest upper bound
4: **if** $\check{f}(\mathcal{S}_{\bar{a}_{\text{parent}}}) - f(\mathcal{S}_{a_{\text{incumbent}}}) < \epsilon$ **then go to** line 11 **end if**
5: **for all** $\bar{a}_{\text{child}}$ from BRANCH($\bar{a}_{\text{parent}}$) **do**
6: **if** $\check{f}(\mathcal{S}_{\bar{a}_{\text{child}}}) > f(\mathcal{S}_{a_{\text{incumbent}}})$ **then**
7: **if** LENGTH($\bar{a}_{\text{child}}$) $= I$ **then**
8: $a_{\text{incumbent}} \leftarrow \bar{a}_{\text{child}}$
9: **else**
10: Append $\bar{a}_{\text{child}}$ to `live`
11: **return** globally optimal $a_{\text{optimal}} = a_{\text{incumbent}}$

1: **function** BRANCH($\bar{a}_{\text{parent}}$)
2: Initialize empty list `children` $= []$
3: **for** $b = 1 : (1 + \max(\bar{a}_{\text{parent}}))$ **do**
4: Append $[\bar{a}_{\text{parent}}, b]$ to `children`
5: **return** `children`
6: **end function**

Proof: First note that $\check{\rho}_{i_k}$ is an upper bound of the achievable long-term SINR for MS i_k in the considered sub-tree, since the requirement of disjoint clusters is not enforced in the optimization problems[8] in (6). We therefore have that $v_{i_k}(|\mathcal{S}_a(i)|, \rho_{i_k}(\mathcal{S}_a(i))) \leq v_{i_k}(|\mathcal{S}_a(i)|, \check{\rho}_{i_k})$, due to the monotonicity property of $v_{i_k}(\cdot, \cdot)$. Now the fact that $v_{i_k}(|\mathcal{S}_a(i)|, \check{\rho}_{i_k}) \leq v_{i_k}(\check{B}_{i_k}, \check{\rho}_{i_k})$ holds is proven. Note that $B_{i_k}^{\star}$ is the optimal size of the cluster, in terms of the first parameter of $v_{i_k}(\cdot, \check{\rho}_{i_k})$. If $|\mathcal{S}_{\bar{a}}(i)| \geq B_{i_k}^{\star}$, the cluster is already larger than what is optimal, and keeping the size is thus a bound for all leaves in the sub-tree. On the other hand, if $|\mathcal{S}_{\bar{a}}(i)| < B_{i_k}^{\star}$ and $i \in \mathcal{P}_{\bar{a}}$, $\check{B}_{i_k}$ is selected as close to $B_{i_k}^{\star}$ as possible, given the number of unconstrained cells that could conceivably be constrained into $\mathcal{S}_{\bar{a}}(i)$ further down in the sub-tree. If $i \in \mathcal{P}_{\bar{a}}^{\perp}$ however, we similarly bound $\check{B}_{i_k}$, except that we only consider cells in non-full clusters for cell i to conceivably be constrained to further down in the sub-tree. Due to the unimodality property of $v_{i_k}(\cdot, \check{\rho}_{i_k})$ and the fact that $\check{B}_{i_k}$ is selected optimistically, we have that $v_{i_k}(|\mathcal{S}_a(i)|, \check{\rho}_{i_k}) \leq v_{i_k}(\check{B}_{i_k}, \check{\rho}_{i_k})$, which gives the bound. ∎

As the algorithm explores nodes deeper in the search tree, LENGTH $(\bar{a})$ gets closer to I, and there is less freedom in the bounds. For LENGTH $(\bar{a}) = I$, the bounds are tight.

2) Algorithm: The proposed branch and bound method is described in Alg. 1. The algorithm starts by getting an initial incumbent solution from a heuristic (e.g. from Sec. III-C or [2]–[4]), and then sequentially studies the sub-tree which currently has the highest upper bound. By comparing the upper bound $\check{f}(\mathcal{S}_{\bar{a}})$ to the currently best lower bound $f(\mathcal{S}_{a_{\text{incumbent}}}) \leq f(\mathcal{S}^{\star})$, the *incumbent solution*, the sub-tree below $\bar{a}$ can be *pruned* if it provably cannot contain the optimal solution, i.e.

[8]The optimal solution to the optimization problem in (6) can be found by minimizing the denominator of $\rho_{i_k}(\mathcal{E}_{i_k})$ in (1), which is easily done using greedy search over the feasible set. The set-function $\rho_{i_k}(\mathcal{E}_{i_k})$ is *supermodular* [12], i.e. demonstrating "increasing returns", which is the structure that admits the simple solution of the optimization problem. Without changes, Thm. 1 would indeed hold for any other supermodular set-function $\rho_{i_k}(\mathcal{E}_{i_k})$.

Algorithm 2 Heuristic for Base Station Clustering

Input: $\mathcal{L} = \{(i,j) \in \mathcal{I} \times \mathcal{I} \mid i \neq j\}$, $\mathcal{S} = \{\{1\}, \ldots, \{I\}\}$

1: **while** $|\mathcal{L}| > 0$ **do**
2: $(i^\star, j^\star) = \arg\max_{(i,j)\in\mathcal{L}} \sum_{k=1}^{K} \log\left(1 + \gamma_{i_k j} P_j / \sigma_{i_k}^2\right)$
3: Let $\Xi_{(i^\star,j^\star)} \leftarrow \mathcal{S}(i^\star) \cup \mathcal{S}(j^\star)$
4: **if** $\left|\Xi_{(i^\star,j^\star)}\right| \leq D$ **then**
5: Let $\mathcal{S} \leftarrow (\mathcal{S} \setminus \{\mathcal{S}(i^\star), \mathcal{S}(j^\star)\}) \cup \{\Xi_{(i^\star,j^\star)}\}$
6: Let $\mathcal{L} \leftarrow \mathcal{L} \setminus \{(i^\star, j^\star)\}$
7: **return** heuristic solution $a_{\text{heuristic}} = a_{\mathcal{S}}$

if $\check{f}(\mathcal{S}_{\bar{a}}) < f(\mathcal{S}_{a_{\text{incumbent}}})$. If a node $\bar{a}$ cannot be pruned, all children of $\bar{a}$ are built by a branching function and stored in a list for future exploration by the algorithm. If large parts of the search tree can be pruned, few nodes need to be explicitly explored, leading to a complexity reduction. The algorithm ends when the optimality gap for the current incumbent solution is less than a pre-defined $\epsilon \geq 0$.

Theorem 2. *Alg. 1 converges to an ϵ-optimal solution of the optimization problem in* (5) *in at most $\sum_{i=1}^{I} \mathbb{B}_i$ iterations.*

Proof: Only sub-trees in which the optimal solution cannot be are pruned. Since all non-pruned leaves are explored, the global optimum will be found. No more than all $\sum_{i=1}^{I} \mathbb{B}_I$ nodes of the search tree can be traversed. ■

In Sec. IV we empirically show that the average complexity is significantly lower than the worst case.

C. Heuristic Base Station Clustering

We also provide a heuristic (see Alg. 2) which can be used as the initial incumbent in Alg. 1, or as a low complexity clustering algorithm in its own right. The heuristic works by greedily maximizing a function of the average channel gains in the clusters while respecting the cluster size constraint. The heuristic is similar to Ward's method [17].

IV. Numerical Results

For the performance evaluation [18], we consider a network of $I = 16$ BSs, $K = 2$ MSs per cell, and $d = 1$ per MS. We employ the throughput model from (4) and let $f(\mathcal{S}) = \sum_{(i,k)} t_{i_k}(\mathcal{S})$. We let the number of antennas be $M = 8$ and $N = 2$. This gives a hard size constraint as $D = 4$ cells per cluster, due to IA feasibility [13]. We consider a large-scale setting with path loss $15.3 + 37.6 \log_{10}(\text{distance [m]})$, i.i.d. log-normal shadow fading with 8 dB std. dev., and i.i.d. $\mathcal{CN}(0,1)$ small-scale fading. The BSs are randomly dropped in a $2000 \times 2000\,\text{m}^2$ square and the BS-MS distance is 250 m. We let $L_c = 2\,700$, corresponding to an MS speed of 30 km/h at a typical carrier frequency and coherence bandwidth [19].

In Fig. 2 we show the convergence of the best upper bound and the incumbent solution, respectively, for one network realization with $\text{SNR} = P_{i_k}/\sigma_{i_k}^2 = 20\,\text{dB}$. The number of iterations needed was 198 and a total of 908 nodes were bounded. Naive exhaustive search would have needed exploring $\mathbb{B}_{16} = 10\,480\,109\,379$ nodes, and the proposed algorithm was thus around $1 \cdot 10^7$ times more efficient for this realization[9]. The number and fraction of nodes pruned during the iterations is shown in Fig. 3. At convergence, 99.99999% of the search tree had been pruned.

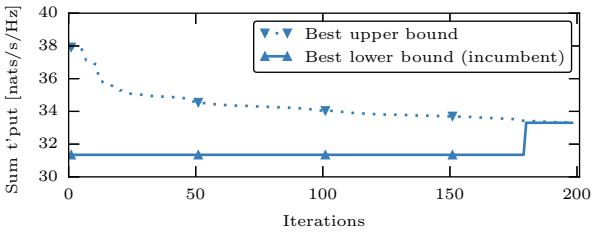

Fig. 2. Example of convergence of the algorithm for one realization.

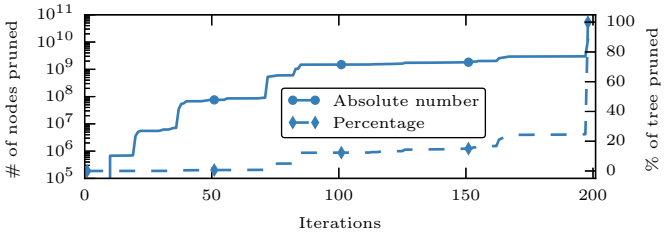

Fig. 3. Pruning evolution for the realization in Fig. 2.

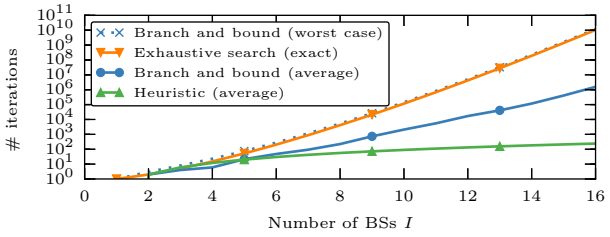

Fig. 4. Average complexity as a function of I.

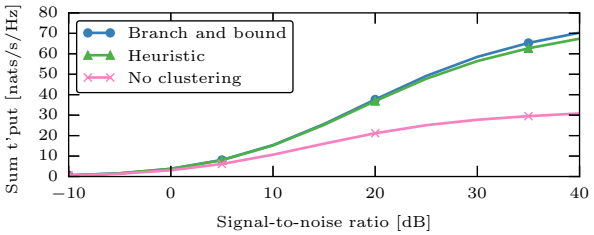

Fig. 5. Sum throughput performance as a function of SNR.

We show the average number of iterations as a function of network size in Fig. 4. The complexity of the proposed algorithm is orders of magnitude lower than the complexity of exhaustive search.

In Fig. 5 we show the sum throughput performance as a function of SNR, averaged over 250 network realizations. The heuristic algorithm performs well: it is close to the optimum[10], and has about twice the throughput of the no clustering case, where $\mathcal{S} = \{\{1\}, \ldots, \{I\}\}$. The grand cluster $\mathcal{S} = \{\mathcal{I}\}$ has zero sum throughput since $I > D$, and is therefore not shown.

[9] Also note that $\sum_{i=1}^{16} \mathbb{B}_i = 12\,086\,679\,035$, i.e. the actual running time of the algorithm was significantly lower than the worst-case running time.

[10] For this network size, exhaustive search is not tractable. Without Alg. 1, we would thus not know that Alg. 2 performs so well. This shows the significance of Alg. 1 as a benchmarking tool for practical but suboptimal clustering algorithms such as Alg. 2, or the algorithms in [2]–[4].

V. Conclusions

With a structured branch and bound approach, the otherwise intractable base station clustering problem has been solved. The algorithm is intended for benchmarking of suboptimal base station clustering heuristics in intermediate size networks.

References

[1] E. Björnson and E. Jorswieck, "Optimal resource allocation in coordinated multi-cell systems," *Foundations and Trends in Communications and Information Theory*, vol. 9, no. 2-3, pp. 113–381, 2013.

[2] S. Peters and R. Heath, "User partitioning for less overhead in MIMO interference channels," *IEEE Trans. Wireless Commun.*, vol. 11, no. 2, pp. 592–603, Feb. 2012.

[3] S. Chen and R. S. Cheng, "Clustering for interference alignment in multiuser interference network," *IEEE Trans. Veh. Technol.*, vol. 63, no. 6, pp. 2613–2624, Jul. 2014.

[4] R. Brandt, R. Mochaourab, and M. Bengtsson, "Distributed long-term base station clustering in cellular networks using coalition formation," *Submitted to IEEE Trans. Signal Inf. Process. Netw.*, 2015, arXiv:1602.07859 [cs.IT].

[5] O. Ayach and R. Heath, "Interference alignment with analog channel state feedback," *IEEE Trans. Wireless Commun.*, vol. 11, no. 2, pp. 626–636, Feb. 2012.

[6] H. Bölcskei and J. Thukral, "Interference alignment with limited feedback," pp. 1759–1763, Jun. 2009.

[7] J. Jose, A. Ashikhmin, T. L. Marzetta, and S. Vishwanath, "Pilot contamination and precoding in multi-cell TDD systems," *IEEE Trans. Wireless Commun.*, vol. 10, no. 8, pp. 2640–2651, 2011.

[8] R. Mochaourab, E. Björnson, and M. Bengtsson, "Adaptive pilot clustering in heterogeneous massive MIMO networks," *Submitted to IEEE Trans. Wireless Commun.*, Jul. 2015, arXiv:1507.04869 [cs.IT].

[9] V. R. Cadambe and S. A. Jafar, "Interference alignment and degrees of freedom of the K-user interference channel," *IEEE Trans. Inf. Theory*, vol. 54, no. 8, pp. 3425–3441, 2008.

[10] A. H. Land and A. G. Doig, "An automatic method of solving discrete programming problems," *Econometrica*, vol. 28, no. 3, pp. 497–520, Jul. 1960.

[11] R. Horst and H. Tuy, *Global Optimization: Deterministic Approaches*. Springer, 1996.

[12] S. T. McCormick, "Submodular function minimization," in *Discrete Optimization*, ser. Handbooks in Operations Research and Management Science, G. N. K. Aardal and R. Weismantel, Eds. Elsevier, 2005, vol. 12, pp. 321–391.

[13] T. Liu and C. Yang, "On the feasibility of linear interference alignment for MIMO interference broadcast channels with constant coefficients," *IEEE Trans. Signal Process.*, vol. 61, no. 9, pp. 2178–2191, May 2013.

[14] D. E. Knuth, *The Art of Computer Programming: Combinatorial Algorithms, Part 1*. Addison-Wesley, 2011, vol. 4A.

[15] R. A. Brualdi, *Introductory Combinatorics*. Pearson, 2009.

[16] D. Berend and T. Tassa, "Improved bounds on Bell numbers and on moments of sums of random variables," *Prob. and Math. Stat.*, vol. 30, no. 2, pp. 185—-205, 2010.

[17] J. Ward, "Hierarchical grouping to optimize an objective function," *J. Amer. Statist. Assoc.*, vol. 58, no. 301, pp. 236–244, Mar. 1963.

[18] R. Brandt, "Julia simulation environment for globally optimal base station clustering in interference alignment-based multicell networks," 2016, available at: https://github.com/rasmusbrandt/OptimalBaseStationClustering.jl.

[19] N. Jindal and A. Lozano, "A unified treatment of optimum pilot overhead in multipath fading channels," *IEEE Trans. Commun.*, vol. 58, no. 10, pp. 2939–2948, Oct. 2010.